GP 2731
#3

Docket No.: 585-1003

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

IN THE APPLICATION OF)
)
Rosenberg et. al.) Group Art Unit No. 2731
)
SERIAL NO: 09/349,347)
)
FILED: July 7, 1999)
)
FOR: Integrated Connection Admission Control and Bandwidth on Demand for Multiple Access ATM-Like Networks)

OIPE JC72
MAR 22 2000
PATENT & TRADEMARK OFFICE

RECEIVED
MAR 24 2000
TC 2700 MAIL ROOM

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to "Commissioner of Patents, Washington, D.C. 20231", on March 17, 2000.

Name of person signing Sheri Fassl

Signature Sheri Fassl

CLAIM FOR PRIORITY

Honorable Commissioner of
Patents and Trademarks
Washington, D.C. 20231

Dear Sir:

Under the International Convention, for the purposes of priority, applicant claims the benefit of United Kingdom Application No. 9815420.6 filed July 16,1998.

A certified copy of said application is appended hereto.

DATE: March 17, 2000

Respectfully submitted,

William M. Lee

William M. Lee, Jr.
Registration No. 26,935

LEE, MANN, SMITH, MCWILLIAMS
SWEENEY & OHLSON
P.O. Box 2786
Chicago, Illinois 60690-2786
(312) 368-1300
Fax (312) 368-0034

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed *A. Evans.*

Dated 15 February 2000

CERTIFIED COPY OF PRIORITY DOCUMENT

9815420.6

The Patent Office

THE PATENT OFFICE
A
16 JUL 1998
RECEIVED BY FAX

16JUL98 E376198-1 D02973
P01/7700 25.00 - 9815420.6

16 JUL 1998

Request for grant of a patent

The Patent Office
Cardiff Road
Newport
Gwent NP9 1RH

1	Your reference	P 60034 GB
2	Patent application number	9815420.6
3	Full name, address and postcode of the applicant	Northern Telecom Limited World Trade Center of Montreal 8th Floor, 380 St Antoine Street West Montreal H2Y 3Y4 Canada
	Patents ADP number	00258966001
	State of incorporation	QUEBEC - CANADA
4	Title of the invention	Integrated Connection Admission Control and Bandwidth on Demand for a Multiple Access ATM Like Network
5	Name of agent	HARRISON GODDARD FOOTE
	Address for service	Vine House, 22 Hollins Lane Marple Bridge Stockport SK6 5BB United Kingdom
	Patents ADP number	00014571003
6	Priority applications	Country / Priority App No / Date of Filing

7	Parent application (eg Divisional)	Earlier Application No	Date of Filing
8	Statement of Inventorship Needed?	Yes	
9	Number of sheets for any of the following (not counting copies of same document)		
	Continuation sheets of this form		
	Description	27	
	Claims	9	
	Abstract		
	Drawings	1	
10	Number of other documents attached		
	Priority documents		
	Translations of priority documents		
	P7/77		
	P9/77		
	P10/77		
	Other documents	**Covering letter, fee sheet**	

11 I/We request the grant of a patent on the basis of this application.

Signature

Harrison Goddard Foote

16 July 1998

12	Name and daytime telephone number of person to contact in the United Kingdom	**Heather Regan** **0161 427 7005**

Rosenberg

Integrated Connection Admission Control and Bandwidth on Demand for a Multiple Access ATM Like Network

FIELD OF THE INVENTION

The present invention relates to the integration of Connection Admission Controllers (CAC) and Bandwidth on Demand Controllers (BoD) on networks which have at least one ATM Like multiple access (MA) network segment. The term ATM Like is used to denote ATM networks and other cell relay systems which guarantee Quality of Service to some types of connections.

In a multiple access (MA) network, ie. a network which contains at least one MA segment, users of a MA segment share a common medium uplink to access the network. For example, users in a MA segment may share a common radio uplink to a satellite so that each user dispersed over a geographical region is able to communicate with the satellite over the link. The common medium uplink may be time divided (TDMA) and so comprise a plurality of time-slots (TS), where a time-slot carries information, also known as traffic, from one user only and the traffic from one user may require more than one time-slot. Alternatively, the common medium uplink may be frequency divided (FDMA) and comprise a plurality of frequency channels. More common is multi-frequency time division (MF-TDMA), in which the common medium uplink is divided into a plurality of frequency channels, each of which is divided into time-slots and traffic from a user is allocated to a certain number of time-slot/frequency channel pairs. A further alternative is code division (CDMA). In this documents the term time-slot (TS) will be used to designate slots from TDMA, FDMA, MF-TDMA, CDMA and other multiple access division techniques.

Examples of MA segments are satellite networks, wireless networks, return path systems for satellites, passive optical networks, adhoc networks and cable networks.

The users in a MA segment communicate with the multiple access network using the common medium uplink via a subscriber access unit (SAU). A SAU may comprise, for example, a terminal which connects an individual user to the network or a concentrator, which connects a group of users (such as a LAN) to the network.

The common medium uplink connects the users in the MA segment with a headend, which is essentially a switch with several inputs (the uplinks) and several outputs (the downlinks) and possibly an interheadend link. The headend will retransmit the traffic towards its destination on the network via a downlink. The downlinks from the headend are generally dedicated links which transport traffic from the headend to destination terminals or concentrators of the MA segment or to other networks via gateways. In the MA segment the downlink link will generally be point-to-multipoint, ie, broadcast, so that each user in the MA segment can receive traffic carried on the downlink. The downlink to a gateway will generally be point-to-point. There may be many headends connected to each other via inter-headend links (IHL).

A network control centre (NCC) controls and manages the communications within the multiple access network. The headend and NCC may be co-located, however, when the headend is a satellite, it is preferred that the headend and NCC are separated with an appropriate division of functionality between the headend and the NCC. This allows the NCC to be ground based because a satellite will have very tight

processing power and buffer space constraints. There may be more than one NCC per MA segment.

Control of ATM-like networks is generally implemented using call or connection admission and policing. Call or connection admission control (CAC) is a network process that receives as input a connection admission request that specifies the traffic descriptor and quality of service (QoS) requirements of the connection and returns a response granting or denying the admission request. The CAC must ensure that the network meets its end-to-end quality of service (QoS) guarantees to connections that are admitted to the network. Therefore, in global networks, CAC is required in both directions of communication for each MA segment and for each point-to-point network over which the requested connection extends. A connection is rejected if at least one of the switches along the network path over which the connection extends does not have the resource to accommodate it.

A further control system, which is oarticularly useful for non-real time connections, required for MA segments only is the bandwidth on demand (BoD) controller which enables a connection to request the resources it requires (eg. extra time slots) on a demand basis while the connection is already in progress in a MA segment where many bursty connections share a common medium uplink.

BoD will generally be invoked many times during the progress of a single connection, while CAC is usually invoked only once, at the connection set-up, for every connection on the network.

Policing is required to ensure that connections do not violate the traffic descriptors declared by the connections during set-up. These traffic descriptors are used by the CAC to determine whether to accept the

connections. Clearly, if the declared traffic descriptors are not adhered to by the connection, QoS is compromised.

It will be clear that CAC, BoD and policing are intricately interconnected in the design of MA networks offering QoS guarantees to connections. The interconnection and interaction between the CAC, BoD and policing have to be carefully designed in order to generate a MA network which is viable and efficient. In particular, it is important that time-slots are used efficiently to enable the resource in the common medium uplink to be used to its best effect so that the number of users that can be supported by the uplink is maximised.

In some MA networks CAC and BoD are performed in the headend, in other networks CAC and BOD are performed in NCCs, or alternatively CAC and BOD are split between headends and NCCs.

OBJECT OF THE INVENTION

The present invention seeks to provide improved integration of CAC and BoD controllers in a MA network which overcomes or at least mitigates one or more of the problems noted above. It is sought to increase the traffic carrying capacity of the common medium whilst complying with the constraints in the headend, particularly when processing and buffer space in the headend is limited because, for example, it is a satellite. It is also sought to couple the CAC for the common medium uplink with the BoD for the common medium uplink in order to provide both QoS to services and efficiency to the system.

SUMMARY OF THE INVENTION

Therefore, according to a first aspect of the present invention there is provided an integrated connection admission control (CAC) and

bandwidth on demand control (BoD) system for allocating the resource of a common medium uplink of a multiple access (MA) network segment, wherein;

the CAC allocates static resource to all virtual channels (VCs) accepted by the CAC,

the CAC books dynamic resource to VCs that require guaranteed dynamic resource, and

the BoD allocates dynamic resource to VCs or to groupings of VCs requesting dynamic resource in such a way that all VCs or groupings of VCs requesting dynamic resource are dynamically allocated at least the guaranteed dynamic resource which has been booked for them by the CAC.

Having at least a proportion of the dynamic resource pre-booked by the CAC enables the common medium uplink to deliver quality of service (QoS) to subscribers because it ensures that each VC will always get the minimum resource it needs. However, by not allocating the booked resource statically and instead allowing the BoD to allocate the booked resource dynamically, any pre-booked dynamic resource that is not needed by a VC or a group of VCs can be reallocated dynamically as extra resource to any other VCs or groups of VCs, within the shared medium uplink, that do need it.

It should be noted that different types of VCs, having, for example, different ATM transfer capabilities will be reserved by the CAC different amounts of static and guaranteed dynamic resource. For example, a CBR (Constant Bit Rate) connection or a rt-VBR (real time-Variable Bit Rate), which are real time connections will generally be allocated static resource only, ie. they will not use BoD. Other VC types, for example, Unspecificed Bit Rate (UBR), will generally be reserved zero static resource, and may also be reserved zero guaranteed dynamic resource,

but instead will have to make do with a share of the remainder or best effort dynamic resource which the BoD is able to allocate to it. nrt-VBR (non real time-Variabel Bit Rate) may be allocated static resource, may be reserved guaranteed dynamic resource and may also rely on a share of the best effort dynamic resource as well. These are system specific examples of how resource would be booked and allocated to a VC and would depend, for example, on propagation delay in the system.

It is preferred, in order to simplify processing in the integrated CAC and BoD system that the CAC reserves static resource for a VC when a VC is set up for the duration of the connection associated with the VC. Similarly, it is preferred that the CAC reserves when applicable booked dynamic resource to a VC when a VC is set up for the duration of the connection associated with the VC.

In order not to exceed the total resource capacity of the medium access uplink the CAC will only accept a VC, designated j, which requires a static resource of SR_j and a booked dynamic resource of BR_j when;

$$SR_j + BR_j + \sum_{k=1}^{K} SR_k + \sum_{k=1}^{K} BR_k \leq C_T$$

where C_T is the total resource capacity of the common medium uplink and there are K existing VCs using the uplink and each of the K VCs have static resource SR_k and booked dynamic resource BR_k reserved to them by the CAC.

It is preferred that the CAC allocates static resource and the BoD allocates dynamic resource on a periodic basis. Furthermore, it is preferred that during a current period the CAC allocates resource for new VCs and releases resource from released VCs for the next period and the BoD allocates dynamic resource for the next period to VCs or

groups of VCs requesting dynamic resource for the next period. As an alternative, the CAC and BoD may be event driven.

In order to simplify processing and to ensure an efficient allocation each period, depending on the status of the network for that period, it is preferred that the allocations made by the BoD for the next period are independent of the allocations made by the BoD for the current period.

It should be noted that the booked dynamic resource for a group of VCs is equal to the sum of the booked dynamic resource reserved for each of the VCs in the group.

Preferably, the BoD allocates dynamic resource to VCs or groups of VCs according to the following rules;

> when the requested resource from the VC or group of VCs is less than or equal to the booked dynamic resource for the VC or group of VCs, the BoD allocates the VC or group of VCs all of the requested resource,
>
> when the requested resource from the VC or group of VCs is greater than the booked dynamic resource for the VC or group of VCs, the BoD allocates the VC or group of VCs the booked dynamic resource and additionally the BoD allocates the VC or group of VCs a share of the remainder of the requested resource, from the remaining resource capacity of the common medium uplink.

This ensures that the BoD complies with the resource bookings made by the CAC and so the BoD will not allocate resource in a way which compromises the Quality of Service guaranteed by the CAC, whilst giving the BoD freedom to allocate the remainder of the resource dynamically so that utilisation of the uplink is optimised. It should be

noted that it is envisaged that a VC or group of VCs may be allocated a zero share of the remainder of the requested resource.

To enable the BoD to further optimise utilisation of the uplink, it is preferred that the share of the remainder of the requested resource for each VC or group of VCs (herein referred to as the best effort resource, BE for each VC or group of VCs), is allocated by maximising the sum of the natural logarithms of all the BEs, subject to the conditions that;

> the BE allocated to each VC or group of VCs is less than or equal to the remainder of the requested resource for that VC or group of VCs, and
>
> the sum of all the BEs is less than or equal to the remaining resource capacity of the common medium uplink.

As an equivalent alternative to maximising the sum of the natural logs of all the BEs, the product of all the BEs can be maximised.

There may be further conditions to this optimisation problem, because in a MF-TDMA common medium uplink, time-slots cannot be allocated to the same SAU simultaneously on different frequency channels. Also, SAUs will have a maximum rate at which they can transmit and receive traffic. The BE allocated to the VCs or groups of VCs originating from the same terminal must take these factors into account.

To guarantee quality of service while providing efficiency for a complete headend comprising many uplinks and many downlinks, further conditions need to be taken into account when maximising the sum of the natural logs of the BEs over the whole set of uplinks. The further conditions will maintain the queues in the buffers of the headend for the downlinks at levels which will prevent loss of traffic cells in the headend. This is especially useful if there is insufficient processing power in the

headend to provide policing in the headend, but is also useful even if there is policing in the headend, because cell loss is reduced. It allows the CAC process to be handled on the basis of static resource and booked dynamic resource only.

Accordingly, where for a complete headend there are multiple uplinks and multiple downlinks controlled by the BoD, it is preferred that the BEs for each uplink are allocated by the BoD by maximising the sum of the log of the BE over the whole set of uplinks subject to the conditions that the sum of all the BEs of a given uplink is less than or equal to the remaining resource capacity of the uplink, for each uplink and the sum of all the BEs for a given downlink is less than or equal to the remaining resource capacity of the downlink, for each downlink and the BE allocated to each VC or group of VCs is less than or equal to the remainder of the requested resource for that VC or group of VCs. The other conditions specific to the terminal would also need to be taken into account if applicable.

According to a second embodiment of the present invention an allocation table setting out resource allocation on the common medium uplink is controlled partly by the CAC when allocating static resource and booking guaranteed dynamic resource and is controlled partly by the BoD when allocating dynamic resource. This provides an efficient way to integrate CAC and BoD activities and share the responsibility of the allocation table in a scalable manner. Thus, it is preferred that the CAC periodically allocates static resource and books guaranteed dynamic resource in the table and then forwards the table to the BoD for the BoD to allocate dynamic resource in the table. It is further preferred that the BoD periodically transmits the completed allocation table, which contains the allocations made by the CAC and the BoD for the next period (possibly shifted by the appropriate propagation delay) to all SAUs.

The CAC is VC based and so fills up the allocation table on a per VC basis and the BoD is generally per group of VC based and so the BoD groups together the relevant VCs before performing the dynamic resource allocation on a per group of VCs basis. For example, the CAC will generally be per connection based and so will fill up the allocation table by allocating or booking each connection to a plurality of time-slots to a connection. Then if the BoD is per terminal based, for example, then, for each terminal, it will group together all the connections from the terminal and give them the same indicator. The booked resource for the VCs from each terminal will be summed to calculate the booked resource for each terminal. Then the BoD will perform the allocation of dynamic resource on a per terminal basis, by allocating each terminal time-slots in the allocation table, from the available dynamic resource. This BoD allocation takes into account each terminal's booked dynamic resource which has been guaranteed by the CAC on a per connection basis.

Note that the CAC books the guaranteed dynamic resource to indicate to the BoD where to find time-slots to use if the resources are requested.

When the headend to which traffic on the uplink is transmitted has low buffer space and processing space, for example, when the headend is a satellite, according to a third aspect of the present invention the CAC and BoD are constrained to allocate resource in such a way that traffic on the common medium access uplink is shaped by the integrated CAC and BoD resource allocation system. The shaping of traffic by the integrated CAC and BoD ensures that the traffic entering the headend complies with the constraints of the ingress of the headend due, in particular, to the low buffer space.

In the above a VC is a connection. The groups of VCs referred to may* have the same source (eg. the same SAU or terminal), may have the same source and have the same ATM transfer capability, may have the same source and the same network destination or may have the same source, the same ATM transfer capability and the same network destination.

The request of resource may be made in terms of a cell rate that will later be transformed by the CAC to a number of allocated time-slots on the common medium uplink. Alternatively, the allocation of resource may be made directly in terms of an allocation of time-slots. It is preferred that the allocation of resource is made in terms of a cell rate, because the computational processing is simplified.

According to a further aspect of the present invention there is provided a method of integrating a connection admission control (CAC) and a bandwidth on demand control (BoD) for allocating the resource of a common medium uplink of a multiple access (MA) network segment, comprising the steps of;

the CAC allocating static resource to all virtual connections (VCs) accepted by the CAC,

the CAC booking dynamic resource to the VCs that require guaranteed dynamic resource, and

the BoD allocating dynamic resource to VCs or to groupings of VCs requesting dynamic resource in such a way that all VCs or groupings of VCs requesting dynamic resource are dynamically allocated at least the guaranteed dynamic resource which has been booked for them by the CAC.

According to a further aspect of the present invention there is provided a method of integrating a connection admission control (CAC) and a

bandwidth on demand control (BoD) system for allocating the resource of a common medium uplink of a multiple access (MA) network segment, which comprises the steps of the CAC allocating static resource and booking guaranteed dynamic resource and the BoD allocating dynamic resource and additionally comprising the steps of filling out an allocation table setting out resource allocation on the common medium access uplink in such a way that the table is controlled by the CAC when allocating static resource and booking dynamic resource and is controlled by the BoD when allocating dynamic resource.

According to a further aspect of the present invention there is provided a method of integrating a connection admission control (CAC) and a bandwidth on demand control (BoD) system for allocating the resource of a common medium uplink of a multiple access (MA) network segment, in which the CAC and BoD are constrained to allocate resource in such a way that traffic on the common medium access uplink is shaped by the integrated CAC and BoD resource allocation system.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will now be described in relation to the accompanying Figures in which;

Figure 1 is a diagram showing the network connections between a SAU, a CAC and a BoD,

Figure 2 shows the periodic timing of the processing conducted by the CAC and the BoD, and

Figure 3 is a schematic representation of part of the CAC allocaton table.

Referring to Figure 2, part of the resource is allocated and booked in an allocation table by a CAC controller periodically, in time periods P. The periodically filled allocation table is sent from the CAC to the BoD controller, periodically as shown by arrows (2). Booked resource as well as left over resource are then allocated in the allocation table by the BoD controller, periodically in periods Q (Q could be equal to P). The BoD will then broadcast the allocation table to the SAUs or terminals, periodically as shown by arrows (4). The allocation table will be effective as soon as it arrives at the terminals, after a propagation delay.

An example of a part of a CAC allocation table is shown in Figure 3. In the CAC table each box represents a time-slot of a MF-TDMA common medium uplink. The different frequencies of the uplink form the columns of the table and the different timings of the uplink form the rows of the table. The CAC has allocated a connection i with three time-slots of static allocation SA_i and reserved the connection i three time-slots of booked dynamic allocation BR_i. Similarly, the CAC has allocated a connection k with three time-slots of static allocation SA_k but has reserved the connection k only one time-slot of booked dynamic allocation BA_k.

In the following the allocation of resource is made in terms of a cell rate, (eg. X cells per second) that will later be transformed by the CAC to a number of allocated time-slots on the common medium uplink.

Applications running on the network will use one of the following ATM transfer capabilities: CBR (Constant Bit Rate); rt-VBR (real time - Variable Bit Rate; nrt-VBR (non-real time - Variable Bit Rate; ABR (Available Bit Rate); GFR (Guaranteed Frame Rate); ABT (ATM Block

Transfer; UBR (Unspecificed Bit Rate) or UBR+. Applications using CBR and rt-VBR are real time applications, the remainder are not. Traffic descriptors are declared by a connection depending on its transfer capabilities.

At the set-up of a connection j, from a terminal, a request for connection is made to the CAC on the common medium uplink. The CAC determines whether it is able to accept the connection (depending on the status of the CAC allocation table as is discussed below) and if it can it will allocate to the connection j a static amount of rate SR_j on the uplink and reserve to the connection a booked rate BR_j (see below). The static rate SR_j is allocated to the connection j for the duration of the connection. The booked rate BR_j is reserved to the connection j for the duration of the connection. When connection j is terminated the CAC will release the rate SR_j and BR_j (see below). It should be noted that the static rate allocated to a connection and the booked rate reserved to a connection could be zero.

The connection j will have an associated traffic descriptor (depending on its ATM transfer capability) and a requested QoS and depending on these characteristics of the connection j, the CAC will allocate an appropriate amount of static rate SR_j. For example, SR_j could be the Peak Cell Rate for a CBR connection. The rate SR_j is chosen by the CAC and it is completely under CAC management.

The connections that are of a transfer capability that do not use BoD, eg. real time connections, will generally only receive static rate SR, but this is system specific and depends on the propagation delay of the system.

For connections that will use the BoD, the CAC will allocate a booked rate BR_j to the connection j which is again dependent on the

characteristics of the connection j. This booked rate could be zero for some connections. The booked rate BR_j will be reserved to connection j for its duration until the connection j is terminated, in which case the booked rate BR_j will be released. Depending on the traffic descriptor and QoS associated with connection j, BR_j could be between zero and PCR_j-SR_j. The following discussion assumes that the BoD is connection based (VC based) and so allocates rate to individual connections j in response to requests from individual connections j.

The allocation of the booked rate, which has been booked by the CAC controller is controlled by the BoD controller. The booked rate BR_j is reserved by the CAC for connection j, so that if connection j requests this booked rate from the BoD controller for the common medium uplink, it is sure to get it. However, if for a period a connection j does not ask for its complete allocation of booked rate, then the remainder of the booked rate is made available for other connections to use on a best effort (BE) basis.

Having this booked rate enables the delivery of QoS to subscriber's by making sure each connection will always get the rate it needs. The advantage of not allocating the booked rate, BR_j statically is that when not needed it can be made available to any other connections within the same uplink.

As indicated above, the CAC on the shared medium uplink will only accept a call j if it has sufficient capacity. That is, if the sum of what is to be statically allocated to connection j (ie. SR_j) and what is to be booked for connection j (ie. BR_j) is less that the total remaining capacity of the uplink, ie. if

$$SR_j + BR_j \leq C_T - \left(\sum_{k=1}^{N} SR_k + \sum_{k=1}^{N} BR_k\right) \qquad (1)$$

where C_T is the total capacity of the shared medium uplink available for traffic and there are N existing connections which have varied static resource SR_k allocated to them and varied booked rate BR_k allocated to them by the CAC where k=1,2,3,....,N. Equation (1) may not be the only condition which the CAC must take into account when deciding to admit a connection because other aspects of the uplink, such as rules for allocating time-slots and aspects of the terminals (in particular the transmit and receive rate of the terminals) may add additional constraints.

A connection j is not restricted to requesting only BR_j. Any connection using the BoD can ask for RR_j from the BoD controller which is less than, equal to or greater than BR_j. It should be noted that the dynamically requested RR_j is in addition to the statically allocated SR_j. When RR_j is greater than BR_j, connection j will be allocated at least rate BR_j by the BoD controller. Connection j will have allocated to it $BR_j + BE_j^*$ where BE_j^* is connection j's fair share of the best effort rate for the next period that the BoD controller will compute knowing the best effort needs $BE_j = \max(0, RR_j - BR_j)$ of all the requests made for that period, according to the following problem:

maximise $\prod_j BE_j^*$ over all connections j of an uplink requesting $RR_j > BR_j$

subject to the constraints $BE_j^* \leq BE_j$

and $\sum_j BE_j^* \leq C_A$

The remaining or best effort rate C_A will be the rate remaining on the uplink after the CAC has allocated all the SR_js, the BoD has allocated all of the RR_i for each connection which does not exceed BR_i and the BOD has allocated that part of the RR_i which does not exceed BR_i for each connection for which RR_i exceeds BR_i.

That is

$$C_A = C_T - \sum_j \min(RR_j, BR_j) - \sum_j SR_j$$

where C_T is the total capacity of the uplink available for traffic.

The allocation of best effort or BE rate is managed completely by the BoD controller.

As discussed above CAC is performed at call set up for all connections and is VC-based (eg. connection based) which means that CAC is performed for each VC. Also, as discussed above BoD can be performed on a VC- basis. However, as discussed below BoD does not have to be performed on a VC basis and, indeed, preferably is not VC based. Alternatively, BoD may be performed as a single operation for a group of VCs (eg. a group of connections) and so does not have to be done on a per VC basis, but may be done on a group of VCs. If BoD is performed on a VC basis it can become unweildly and difficult to implement because of the large number of messages which have to be sent to and from the BoD controller, ie. one set of messages per VC. When BoD is not VC based there is a further problem when integrating CAC and BoD functions because CAC is inherently VC based. Also the CAC and BoD have to have some means of mapping the VC based allocations and bookings made by the CAC to the group of VC (eg. terminal) based allocations made by the BoD and vice versa. This mapping can be set when the system is designed or may be dynamic as a result of communications between the CAC and the BoD.

In what follows, the BoD is not conducted on a VC basis, but on a terminal basis, ie. the BoD allocates rate, relating to all the connections using BoD running on that terminal, to that terminal in a single operation

in response to a request from the terminal, which request groups together the requirements of all the connections. When the terminal has more than one connection running it will itself allocate time-slots to a particular connection from the time-slots allocated to the terminal by the BoD controller. Thus, for a given period the BoD will give a terminal t requesting resource RR_t, either

$$RR_t \qquad \text{if } RR_t \leq \sum_{i=1}^{M} BR_i$$

or

$$\sum_{i=1}^{M} BR_i + BE_t^{*} \qquad \text{if } RR_t > \sum_{i=1}^{M} BR_i \qquad (2)$$

where there are M connections i (i = 1,2,3......M) per terminal, and BE_t^{*} is the terminal's fair share of the best effort rate available for the next period, which the BoD will compute knowing the best effort requirements of all terminals.

If the headend has sufficient processing and buffer space then it can itself shape traffic at the ingress to the headend to conform to the requirements for traffic at the ingress of the headend. In this case the CAC and BoD on the common medium uplink does not need to be strongly coupled with the process in the headend. In other words traffic entering the headend can be as bursty as necessary to achieve maximum utilisation of the common medium uplink, because the headend has sufficient processing and buffer space to transform the incoming bursty traffic into traffic which conforms to the requirements at the ingress of the headend. This means that the CAC and BoD for the uplink can be designed with only maximum utilisation of the common medium uplink in mind. This case is discussed below. Thus the rate available for best-effort in the common medium uplink can be shared among current best effort requests without worrying about the impact of this sharing on the congestion in the buffers of the headend.

At the beginning of each period the BoD controller has to share among terminals requesting best effort rate for the next period, the available capacity C_A which is the rate remaining on the uplink after;

1. the SR_i has been allocated to each connection continuing into or starting in the next period by the CAC,

2. all the RR_t that does not exceed booked rate have been allocated by the BoD controller, and

3. all the RR_t that exceed the booked rate have had allocated to it by the BoD the booked rate $\sum_{i=1}^{n} BR_i$ reserved to it by the CAC.

That is

$$C_A = C_T - \sum_{t} \min(RR_t, BR_t) - \sum_{t} SR_t$$

$$\text{When } SR_t = \sum_{i \in t} SR_i \text{ and } BR_t = \sum_{i \in t} BR_i$$

The BoD has to share the remaining available capacity C_A between all the terminals t requesting best effort resource $BE_t = \max(0, RR_t - BR_t)$. The BoD calculates for each period for each terminal t having a non zero BE_t (assuming there are N such terminals in the period) ($t = 1,2,3....,N$) a best effort allocation BE_t^* by solving the following problem;

maximising $\prod_{t=1}^{N} BE_t$ over all t
subject to the constraint that $BE_t^* \leq BE_t$
and that $\sum_{t=1}^{N} BE_t^* \leq C_A$.

Alternatively, maximise the sum of the logs of BE_t.

This optimisation problem may have additional constraints depending on other aspects of the network or the terminal.

In this example the connections have been grouped, for the BoD process, as connections from the same terminal that require BoD. It should be noted this grouping could be subdivided, for example into a sub-set of connections from the same terminal that require BoD and that additionally have the same destination, or that additionally have the same ATM transfer capabilities or both.

In general where the common medium uplink is divided on a MF-TDMA basis, ie. into time-slots and frequency channels, most terminals have the constraint that they cannot send information on more than one frequency channel simultaneously. In order to overcome this constraint the BoD controller computes BE_t^* using the above optimisation procedure and then translates BE_t^* into a number of time-slots. Then the BoD tries to allocate the time-slots to the terminals so that the MF-TDMA constraints on terminals are complied with. If, due to these constraints, some terminals cannot use their full allocation of time-slots, then the BoD can allocate them to other terminals.

A burst time plan (BTP), is an allocation plan or allocation table (AT) for the next period which is generated, for each period, by the CAC and BoD to show which VCs and/or groups of VCs are allocated which time-slots. To do this the CAC and BoD share the responsibility for preparing the allocation table and have to exchange information in order to generate the BTP. The CAC allocation table or CAT, is filled up by the CAC and then sent to the BoD which uses the CAT to create the allocation table (ie. the BTP). In creating the AT, the BoD controller uses the CAT that the CAC sends to it periodically.

Every period, the BoD controller broadcasts on the downlink the full burst time plan for the corresponding uplink. The BTP will contain an ordered array of time-slot fields, with each field containing the identity of the BoD groups of VCs, if the time-slot has been allocated by the BoD or the identity of the individual VCs if the time-slot has been statically allocated by the CAC.

For new connections arriving in the current period P based of the rate SR, the CAC computes the number of statically allocated time-slots (SATS) that must be allocated to the new VCs. Similarly, based on the rate BR, the CAC computes the number of booked time-slots (BATS) to be reserved for the new connection. Next, the CAC updates the CAT with the newly allocated and booked time-slots after releasing the time-slots allocated and booked for connections that have just been released. This table is sent to the BoD controller periodically. Alternatively the CAT can be sent to the BoD controller on an event driven basis.

Hence, the CAC can use the CAT to deallocate SATS and BATS for a connection that has been terminated and to allocate SATS and BATS for new connections. All these activities take effect only when the corresponding AT created by the BoD from the CAT are received by the terminals.

Thus, the CAT comprises at the beginning of a given period, the current view of the uplink CAC of the status of the time-slots in the next period. Each time-slot (TS) in the CAT can be unused, SATS or BATS. The CAT time-slots that are SATS have been assigned by the CAC for the next period and the BATS have been pre-allocated but not allocated. Each SATS or BATS has an identifier for VC i that owns it. The CAC works on VC, ie. on connection and so this implies that the CAT is VC-based. In this case the BoD can make a transformation replacing the

VC identifier of the BATS (but not the SATS) with the identifier of the group of VCs to which the VC belongs.

A new connection is accepted by the uplink CAC only if the connection's required SATS and BATS can be assigned to currently unused time-slots in the CAT while respecting constraints on the SAUs, eg. terminal constraints. Note that where terminal constraints are not complied with, for example, because there is no way not to allocate two time slots simultaneously on different channels to the same terminal while there is in principle enough time slots to accept the connection, the CAT may be rearranged to move some of the allocated time-slots to a different time. Hence, the CAT may be re-arranged from time to time.

Designing the CAC and BoD so that they each perform separate defined functions enables the CAC controller and the BoD controller to be located in different locations on the network or enable improved scalability when they are co-located. In the case of a satellite network, both controllers could be co-located in a ground based NCC, or the CAC controller could be located in the ground based NCC with the BoD controller located in the satellite.

Once the overall CAC decides to grant admission to connection j, a virtual path identifier (VPI) and a virtual channel identifier (VCI) are assigned to connection j and the SAU is informed of this VPI/VCI and that connection j has been granted admission to the network. The CAC also informs the BoD controller by communicating to it the identifiers of all the new connections (and possibly all the terminated connections) once each period at the same time as it sends the CAT. When the BoD controller receives at the beginning of a period the identifiers for all newly set-up connections j, possibly the identifiers of the connections that have been terminated and the CAT, the BoD controller can enter

into a BoD response phase, in which it allocates all the dynamic rate taking into account the SATS and the BATS in the CAT.

Figure 3 illustrates the coupling between the CAC and BoD controllers. The periodicity of the CAC/BoD process should be chosen to minimise overhead, to deliver a good efficiency and to allow acceptable timing at connection set-up and acceptable reactiveness of the BoD Protocol.

A stream of traffic passing through a multiple access uplink will almost certainly be transformed, in terms of its traffic descriptor by this passage.

If there are few constraints on the buffer space and processing space in the headend then the CAC in the headend and the CAC/BoD for the uplink can be kept separate so that CAC/BoD receiving traffic having traffic descriptor X, allocates time-slots on the uplink in the most optimised manner as discussed above. The consequence of this may be to transform the traffic into traffic which could be less smooth than the original one and having traffic descriptor L. Then in order not to create any problem in the headend, the headend will have to re-shape traffic having descriptor L at its ingress to conform to the traffic descriptor X.

If the headend is constrained, eg. as in the case of a satellite, then reshaping cannot be performed in the headend, so that the CAC/BoD for the uplink, at the cost perhaps of full optimisation of time-slot allocation, is designed to create a cell stream, with a traffic descriptor that can be accepted by the headend. If a connection of traffic descriptor X is very bursty and the headend has very low buffer space, there is a high probability that it will not be accepted by the CAC in the headend. But particularly if this traffic is not real time, it could be transformed into traffic having a descriptor K_X, which is smoother and

less bursty that would allow the CAC in the headend to accept it with a higher probability. Thus, the CAC decision in the headend can be based on K_X as long as the CAC/BoD on the uplink can transform connections having traffic descriptor X into traffic conforming to traffic descriptor K_X. If the thus transformed traffic goes through a gateway to another part of the network, the gateway might have to transform the traffic again so that it conforms once more with the original traffic descriptor X. If the traffic goes directly to another SAU, then that SAU has to be informed that it will receive something that corresponds to K_X. Note that the transformation $X \Rightarrow K_X$ and the subsequent transformation $K_X \Rightarrow X$ will involve delaying some cells.

If the traffic with descriptor X is real-time, then the CAC/BoD on the uplink may only be able to accept the traffic by transforming it into traffic with a descriptor K_X which is less smooth than X.

Three examples of such traffic shaping by the CAC and BoD are given below:

1. If connection i is CBR (Constant Bit Rate), which is real-time, then the BoD controller may not be used if propagation delay is large. If the delay constraints for i are very tight, then the uplink CAC/BoD might have to transform connection i into a CBR of PCR (Peak Cell Rate) greater than the one declared by connecton i.

2. If connection i is rt-VBR (real time variable bit rate) then the uplink CAC will transform connection i having traffic descriptor X into a connection j having traffic descriptor K_X not conforming to the traffic descriptor X. Connection i could be transformed into a CBR of $PCR_j > SCR_i$. The headend CAC has then to accept i based on PCR_j.

3. If connection i is nrt-VBR, then the CAC/BoD could transform connection i without effecting its QoS. Then this should be done in a way that helps the CAC in the headend. An extreme example would be to transform connection i into a CBR_j of $PCR_j = SCR_i$ (eg.$SR_i = SCR_i$).

This will not by itself avoid congestion in the headend, since it does not take into account the allocation of best effort time-slots. If the BoD controller allocates time-slots without taking into account the state of the buffers in the headend, there remains the possibility of congestion which at the very best would then mean the discarding of best effort cells that have been sent on the shared uplink at a high cost. Thus, to avoid congestion the SAUs can make sure that each connection conforms to the traffic requirements of the CAC in the headend. However, this does not solve the problem for connections which rely mostly on best effort, as best effort traffic (ie. on top of booked rate) is not subject to the CAC process.

To solve this problem the computation of the fair share of best effort rate in one common medium uplink is co-ordinated with all the other uplinks arriving at the headend as well as all the downlinks from the headend (assuming there is only one headend).

In order to co-ordinate the BoD on all the uplinks of a given headend with all its downlinks to avoid congestion when the buffer space for each downlink is very limited, rate allocation is performed as discussed below. This will allow the CAC process in the headend to be performed based only on static rate and booked rate in the uplink since all the best effort resource will be well regulated.

A request for a best effort resource, ie. a best effort rate, has to be done on behalf of a group of VCs from the same terminal (operating on an uplink) destined for the same downlink or a subset of this grouping of

VCs. That is requests for rate have to be made in such a way that the BoD controller knows for which downlink the request is for. This grouping of VCs is referred to below as a user.

Let the headend have M uplinks and M downlinks where u_i is uplink i where $0<i<M+1$ and d_j is the downlink j $(0<j<M+1)$, let x be the best effort rate that should be allocated to a user on top of its aggregated static rate SR and booked rate BR for the next period knowing the requested best effort rate BE for the next period of each of the users in all the uplinks and downlinks. If the BoD controller is centralised for all the uplinks and downlinks and if there is only one headend, the best effort rates x can be computed as follows.

In the current period, there are n_{ij} requests made for the next period for d_j from u_i. Let BE_t^{ij} $(1 \leq t \leq n_{ij})$ be the best effort request from best-effort user t of uplink i for downlink j (for $1 \leq i,j \leq M$). Then the capacity available for best effort within each uplink i (C_{Ai}) and the capacity available for best effort within each downlink (D_{Aj}) has to be shared in an efficient and fair way between the users. C_{Ai} can be computed as detailed above for C_A for each uplink for the next period to the headend. The capacity available for a given downlink j for the next period is:

$$D_{Aj} = C_{Dj} - \sum_{i=1}^{M} \sum_{t=1}^{n_{ij}} \min(RR_t^{ij}, \sum_{v \in entity\, t} BR_v\, \Im(V \in d_j)) - \sum_{i=1}^{M} \sum_{k} SR_k\, \Im(k \in u_i)\Im(k \in d_j)$$

Where C_{Dj} is the total capacity of d_j and $\Im(\,)$ is the indicator function, ie. $\Im(x) = 1$ if x is true and otherwise zero.

Then the allocation of rate x_k^{ij} to user k in uplink i, for downlink j is calculated as follows:

$$\text{Maximise} \sum_{i=1}^{M} \sum_{j=1}^{M} \sum_{k=1}^{n_{ij}} Ln(X_k^{ij}) \text{ subject to}$$

$x_k^{ij} \leq BE_k^{ij}$

$\sum_{i=1}^{m} \sum_{k=1}^{n_{ij}} x_k^{ij} \leq D_{Aj}$ for all the downlinks j

$\sum_{j=1}^{m} \sum_{k=1}^{n_{ij}} x_k^{ij} \leq C_{Aj}$ for all the uplinks i, and

$\sum_{j=1}^{m} \sum_{k=1}^{n_{ij}} x_k^{ij} \leq M_{termp}$ for all terminals p in the uplinks i where M_{termp} is the maximum traffic rate of terminal p.

CLAIMS

1. An integrated connection admission control (CAC) and bandwidth on demand control (BoD) system for allocating the resource of a common medium uplink of a multiple access (MA) network segment, wherein;

the CAC allocates static resource to all virtual connections (VCs) accepted by the CAC,

the CAC books dynamic resource to the VCs that require guaranteed dynamic resource, and

the BoD allocates dynamic resource to VCs or to groupings of VCs requesting dynamic resource in such a way that all VCs or groupings of VCs requesting dynamic resource are dynamically allocated at least the guaranteed dynamic resource which has been booked for them by the CAC.

2. A system according to claim 1 wherein the CAC allocates static resource to a VC when a VC is set up for the duration of the connection associated with the VC.

3. A system according to claim 1 or claim 2 wherein the CAC reserves booked dynamic resource to a VC when a VC is set up for the duration of the connection associated with the VC.

4. A system according to any one of the proceeding claims wherein the CAC will only accept a VC j requiring a static resource of SR_j and a booked dynamic resource of BR_j when;

$$SR_j + BR_j + \sum_{k=1}^{K} SR_k + \sum_{k=1}^{K} BR_k \leq C_T$$

where C_T is the total resource capacity of the common medium uplink and there are K existing VCs using the uplink and each of the K VCs

have been reserved static resource SR_k and booked dynamic resource BR_k by the CAC.

5. A system according to any one of the preceding claims wherein the CAC allocates reserved static resource and the BoD allocates dynamic resource on a periodic basis.

6. A system according to claim 5 wherein during a current period the CAC allocates resource for new VCs and deallocates resource from released VCs for the next period and the BoD allocates dynamic resource for the next period to VCs or groups of VCs requesting dynamic resource for the next period.

7. A system according to claim 5 or claim 6 wherein the allocations made by the BoD for the next period are independent of the allocations made by the BoD for the current period.

8. A system according to any one of the preceding claims wherein the BoD allocates dynamic resource to VCs or groups of VCs according to the following rules;

when the requested resource from the VC or group of VCs is less than or equal to the booked dynamic resource for the VC or group of VCs, the BoD allocates the VC or group of VCs all of the requested resource,

when the requested resource from the VC or group of VCs is greater than the booked dynamic resource for the VC or group of VCs, the BoD allocates the VC or group of VCs the booked dynamic resource and additionally the BoD allocates the VC or group of VCs a share of the remainder of the requested resource, from the remaining resource capacity of the common medium uplink.

9. A system according to claim 8 wherein the share of the remainder of the requested resource, or the best effort resource, BE, for each VC or group of VCs, is allocated by maximising the sum of the natural logarithms of all the BEs, subject to the conditions that;

the BE allocated to each VC or group of VCs is less than or equal to the remainder of the requested resource for that VC or group of VCs, and

the sum of all the allocated BEs is less than or equal to the remaining resource capacity of the common medium uplink.

10. A system according to claim 8 wherein the share of the remainder of the requested resource, or the best effort resource, BE, for each VC or group of VCs, is allocated by maximising the product of all the BEs, subject to the conditions that;

the BE allocated to each VC or group of VCs is less than or equal to the remainder of the requested resource for that VC or group of VCs, and

the sum of all the allocated BEs is less than or equal to the remaining resource capacity of the common medium uplink.

11. A system according to claim 9 wherein there are multiple uplinks and multiple downlinks controlled by the BoD for a single headend, the BEs for each uplink are allocated by the BoD subject to the additional conditions that;

the sum of all the allocated BEs is less than or equal to the remaining resource capacity of the uplink, for each uplink, and

the sum of all the allocated BEs is less than or equal to the remaining resource capacity of the downlink, for each downlink.

12. A system according to any one of the preceding claims which has an allocation table setting out resource allocation on the common medium access uplink which is controlled by the CAC when allocating static resource and booking dynamic resource and is controlled by the BoD when allocating dynamic resource.

13. A system according to claim 12 wherein the CAC periodically allocates resource and books resource in the table and forwards the table to the BoD to allocate dynamic resource in the table.

14. A system according to any one of the preceding claims wherein the CAC and BoD are constrained to allocate resource in such a way that traffic on the common medium access uplink is shaped by the integrated CAC and BoD resource allocation system.

15. An integrated connection admission control (CAC) and bandwidth on demand control (BoD) system for allocating the resource of a common medium uplink of a multiple access (MA) network segment, wherein the CAC allocates static resource and books guaranteed dynamic resource and the BoD allocates dynamic resource and the system comprises an allocation table setting out resource allocation on the common medium access uplink which is controlled by the CAC when allocating static resource and booking dynamic resource and is controlled by the BoD when allocating dynamic resource.

16. A system according to claim 15 wherein the CAC periodically allocates resource and books resource in the table and forwards the table to the BoD to allocate dynamic resource in the table.

17. An integrated connection admission control (CAC) and bandwidth on demand control (BoD) system for allocating the resource of a

common medium uplink of a multiple access (MA) network segment, wherein the CAC and BoD are constrained to allocate resource in such a way that traffic on the common medium access uplink is shaped by the integrated CAC and BoD resource allocation system.

18. A method of integrating a connection admission control (CAC) and a bandwidth on demand control (BoD) for allocating the resource of a common medium uplink of a multiple access (MA) network segment, comprising the steps of;

the CAC allocating static resource to all virtual connections (VCs) accepted by the CAC,

the CAC booking dynamic resource to the VCs that require guaranteed dynamic resource, and

the BoD allocating dynamic resource to VCs or to groupings of VCs requesting dynamic resource in such a way that all VCs or groupings of VCs requesting dynamic resource are dynamically allocated at least the guaranteed dynamic resource which has been booked for them by the CAC.

19. A method according to claim 18 in which the CAC allocates static resource to a VC when a VC is set up for the duration of the connection associated with the VC.

20. A method according to claim 18 or claim 19 in which the CAC reserves booked dynamic resource to a VC when a VC is set up for the duration of the connection associated with the VC.

21. A method according to any one of claims 18 to 20 in which the CAC accepts a VC j requiring a static resource of SR_j and a booked dynamic resource of BR_j only if;

$$SR_j + BR_j + \sum_{k=1}^{K} SR_k + \sum_{k=1}^{K} BR_k \leq C_T$$

where C_T is the total resource capacity of the common medium uplink and there are K existing VCs using the uplink and each of the K VCs have been reserved static resource SR_k and booked dynamic resource BR_k by the CAC.

22. A method according to any one of claims 18 to 21 in which the CAC allocates reserved static resource and the BoD allocates dynamic resource on a periodic basis.

23. A method according to claim 22 in which during a current period the CAC allocates resource for new VCs and deallocates resource from released VCs for the next period and the BoD allocates dynamic resource for the next period to VCs or groups of VCs requesting dynamic resource for the next period.

24. A method according to claim 22 or claim 23 in which the allocations made by the BoD for the next period are independent of the allocations made by the BoD for the current period.

25. A method according to any one of claims 18 to 24 in which the BoD allocates dynamic resource to VCs or groups of VCs according to the following rules;

when the requested resource from the VC or group of VCs is less than or equal to the booked dynamic resource for the VC or group of VCs, the BoD allocates the VC or group of VCs all of the requested resource,

when the requested resource from the VC or group of VCs is greater than the booked dynamic resource for the VC or group of VCs, the BoD allocates the VC or group of VCs the booked dynamic resource and additionally the BoD

allocates the VC or group of VCs a share of the remainder of the requested resource, from the remaining resource capacity of the common medium uplink.

26. A method according to claim 25 in which the share of the remainder of the requested resource, or the best effort resource, BE, for each VC or group of VCs, is allocated by maximising the sum of the natural logarithms of all the BEs, subject to the conditions that;

the BE allocated to each VC or group of VCs is less than or equal to the remainder of the requested resource for that VC or group of VCs, and

the sum of all the allocated BEs is less than or equal to the remaining resource capacity of the common medium uplink.

27. A method according to claim 25 in which the share of the remainder of the requested resource, or the best effort resource, BE, for each VC or group of VCs, is allocated by maximising the product of all the BEs, subject to the conditions that;

the BE allocated to each VC or group of VCs is less than or equal to the remainder of the requested resource for that VC or group of VCs, and

the sum of all the allocated BEs is less than or equal to the remaining resource capacity of the common medium uplink.

28. A method according to claim 26 in which there are multiple uplinks and multiple downlinks controlled by the BoD for a single headend, the BEs for each uplink are allocated by the BoD subject to the additional conditions that;

the sum of all the allocated BEs is less than or equal to the remaining resource capacity of the uplink, for each uplink, and

the sum of all the allocated BEs is less than or equal to the remaining resource capacity of the downlink, for each downlink.

29. A method according to any one of claims 18 to 29 in which an allocation table setting out resource allocation on the common medium access uplink is controlled by the CAC when allocating static resource and booking dynamic resource and is controlled by the BoD when allocating dynamic resource.

30. A method according to claim 29 in which the CAC periodically allocates resource and books resource in the table and forwards the table to the BoD to allocate dynamic resource in the table.

31. A method according to any one of claims 18 to 30 wherein the CAC and BoD are constrained to allocate resource in such a way that traffic on the common medium access uplink is shaped by the integrated CAC and BoD resource allocation system.

31. A method of integrating a connection admission control (CAC) and a bandwidth on demand control (BoD) system for allocating the resource of a common medium uplink of a multiple access (MA) network segment, which comprises the steps of the CAC allocating static resource and booking guaranteed dynamic resource and the BoD allocating dynamic resource and additionally comprising the steps of filling out an allocation table setting out resource allocation on the common medium access uplink in such a way that the table is controlled by the CAC when allocating static resource and booking dynamic

resource and is controlled by the BoD when allocating dynamic resource.

32. A method according to claim 31 in which the CAC periodically allocates resource and books resource in the table and forwards the table to the BoD to allocate dynamic resource in the table.

33. A method of integrating a connection admission control (CAC) and a bandwidth on demand control (BoD) system for allocating the resource of a common medium uplink of a multiple access (MA) network segment, in which the CAC and BoD are constrained to allocate resource in such a way that traffic on the common medium access uplink is shaped by the integrated CAC and BoD resource allocation system.

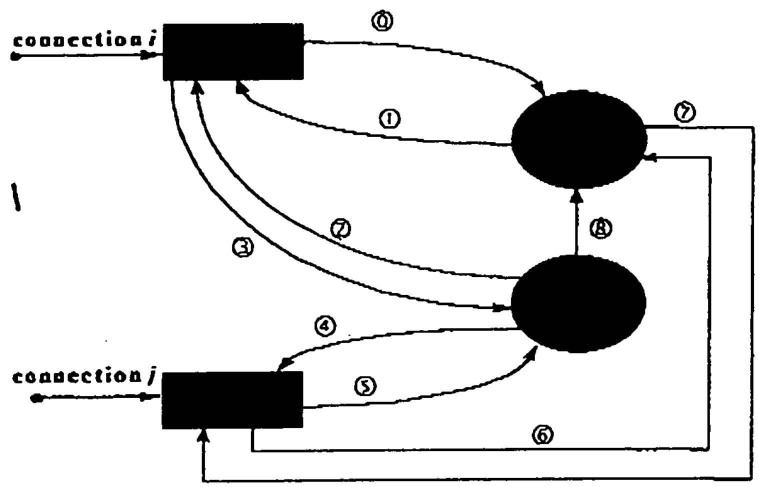

FIG 1

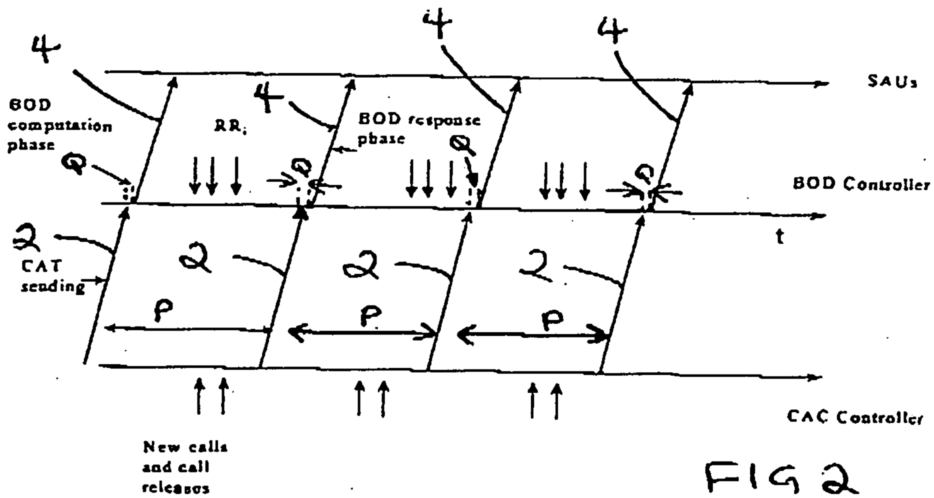

FIG 2

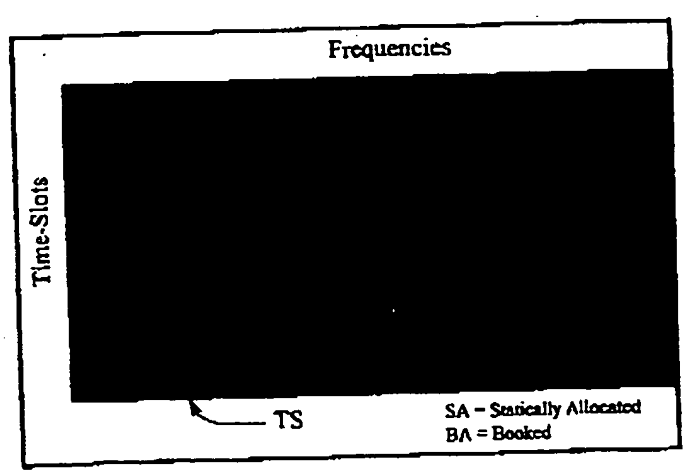

FIG 3.